



Intel® Technology Journal

Tera-scale Computing

Package Technology to Address the Memory Bandwidth Challenge for Tera-scale Computing

Package Technology to Address the Memory Bandwidth Challenge for Tera-scale Computing

Lesley Anne Polka, Assembly and Test Technology Development Division, Intel Corporation
Huthasana Kalyanam, Assembly and Test Technology Development Division, Intel Corporation
Grace Hu, Assembly and Test Technology Development Division, Intel Corporation
Satish Krishnamoorthy, Assembly and Test Technology Development Division, Intel Corporation

Index words: package technology, memory bandwidth, package design, electrical performance

ABSTRACT

Tera-scale computing stresses the platform architecture with memory bandwidth being a likely bottleneck to processor performance that presents unique challenges to CPU packaging. This paper describes the *evolution* in packaging technology with each processor generation to meet increasing memory bandwidth needs and the *revolution* in package technology required for tera-scale computing needs. The scope and focus of the paper are primarily design and electrical performance challenges. We discuss a potential roadmap of transitions in package architecture and technology that evolves from today's off-package memory scenario to increasingly complex on-package integrated memory architectures. An overall treatment of memory hierarchy, including off-die memory approaches, is not within the scope of this paper, but relevant to the overall challenge of enabling higher bandwidth. Again, the focus of this paper is on the CPU package itself. In this context, we discuss the memory bandwidth limitations, technology challenges, and tradeoffs of each package architecture.

INTRODUCTION

With a potential transition to tera-scale computing with multi- and many-core microprocessors and integrated memory controllers on the CPU, memory bandwidth becomes a bottleneck to processor performance [1]. This presents unique challenges to CPU packaging. Previous memory bandwidth requirements have scaled steadily, but fairly slowly, from one microprocessor generation to the next. This has driven a fairly steady but slow increase in pin count growth for chipset packages, which have traditionally provided the link to system memory between the microprocessor and memory modules. With a transition to multi- and many-core architectures, however, there is a large increase in the memory bandwidth requirement. This transition occurs at the same time as a

shift to an integrated memory controller architecture for the CPU. These fairly simultaneous architecture transitions result in a tremendous burden on CPU packaging requirements, driving pin count growth and driving up routing density due to the large increase in interconnects that must be routed from the CPU through the package to off-package memory modules.

In this paper we describe the evolution in packaging technology with each processor generation to meet increasing memory bandwidth needs. We focus on the revolution in package technology required for tera-scale computing needs. The scope and focus of this paper are primarily design and electrical performance challenges. We propose a roadmap of transitions in package architecture and technology that evolves from today's off-package memory to increasingly complex on-package integrated memory architectures. We discuss the memory bandwidth limitations, technology challenges, and tradeoffs of each package architecture.

In the first section of this paper we look at memory bandwidth fundamentals. Next, we review the past trends in memory bandwidth requirements and the package technology impact. We follow this with sections describing the memory bandwidth needs for tera-scale computing and the resulting package technology impact and response.

MEMORY BANDWIDTH FUNDAMENTALS

It is useful to review several fundamental concepts as an introduction to the topic of memory bandwidth. First, it is important to understand the definition of memory bandwidth, the key elements related to bandwidth, and the role that the package interconnect plays. Very basically, memory bandwidth is defined as the product of the

number of data bits in the memory bus and the speed of a single bit in the bus. This can be expressed as

$$\text{BW} = \# \text{ of bits} \times \text{bit rate} \quad \text{Eq. (1)}$$

For example, if a memory bus is 8 bits wide (or 1 byte wide) and each bit transmits data at 1Gb/s (gigabits per second), then the memory bandwidth is 1 byte (1B) x 1Gb/s, or 1GB/s. A more realistic example is that of a typical DDR2 bus that is 16 bytes (128 bits) wide and operating at 800Mb/s. The memory bandwidth of that bus is 16 bytes x 800Mb/s, which is 12.8GB/s.

Besides the actual memory bandwidth, other key elements of memory bandwidth are latency and capacity. Latency is the roundtrip time that it takes to receive a response after a request has been sent. Latency is typically measured in nanoseconds (ns). Capacity refers to the size of the memory and is typically measured in MBs.

The memory subsystem hierarchy of a computer architecture consists of many levels. Memory can be located at the chip level, the package level, the board level, and in separate devices off the board (such as the hard disk). There is a tradeoff among the types and the key elements of memory (bandwidth, latency, and capacity) depending upon the location in the memory subsystem hierarchy. Very simply, faster, lower capacity memory is typically located on-chip, while slower, higher capacity memory is located off-chip. On-chip memory usually uses Static Random Access Memory (SRAM) technology, which is fast but expensive, and it is low-density compared to other memory technologies. On-chip memory usually serves as a cache and can be further divided into levels of cache, e.g., L1 cache, L2 cache, etc., [2]. Off-chip memory typically uses Dynamic Random Access Memory (DRAM) technology, which is slower but cheaper, and it is higher-density than SRAM. Off-chip memory located on the system board serves as the main memory for the computer system.

Today's typical computer architecture consists of the microprocessor (CPU), the chipset, and the main memory. Busses connect the various components of the system. Figure 1 illustrates a typical system architecture consisting of a microprocessor connected to a chipset through the system bus. The chipset in this example is divided into a Memory Controller Hub (MCH) and a separate Graphics Processing Unit (GPU). Each has a memory bus connecting to on-board memory. The system bus connects the CPU to the on-board, main system memory.

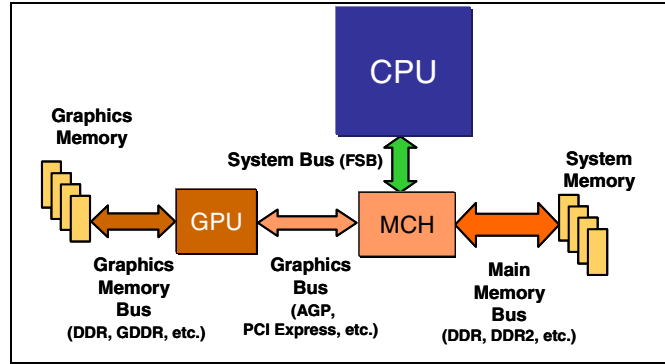


Figure 1: System architecture with main memory bus connected to the chipset

In this example, there are potential bottlenecks at each interconnect transition with respect to providing memory bandwidth to the CPU. Specifically, there is a transition from the CPU to the MCH through the system bus; and there is a transition from the MCH to the system memory through the main memory bus. The challenges to memory bandwidth in this traditional architecture have been to increasingly scale the capabilities of both the system and main memory busses to keep up with the steadily increasing memory demand of the CPU with each new generation. Figure 2 illustrates the historical trending of the system bus bandwidth capability vs. the memory bandwidth of the system. It makes sense that the two bandwidths have needed to scale simultaneously for optimum system performance. Scaling of bus capability has usually involved a combination of increasing the bus width while simultaneously increasing the bus speed.




Figure 2: Historical trend for system bus bandwidth vs. memory bandwidth

Inherent in achieving both increased bus width and increased bus speed is the ability to scale the package technology to accommodate wider and faster busses. The package is the transitional interconnect between the fine pitch and high-density features of the processor chip and

the much coarser pitch and low-density features of the system board. The ability of packaging technology to serve as an intermediary interconnect bridging the gap between the chip and the system board has been critical to enabling increasing system memory bandwidth in the past. Packaging technology will continue to play a critical and increasingly important enabling role as we transition to tera-scale computing architectures.

REVIEW OF PACKAGE TECHNOLOGY EVOLUTION VS. MEMORY BANDWIDTH REQUIREMENTS

In the traditional system architecture of Figure 1, the packaging challenges associated with ever-evolving and increasing memory bandwidth impacted both chipset and CPU packaging. The chipset package has usually absorbed the need for increasing numbers of connections to system memory, while the CPU packaging dealt primarily with the need for an interconnect that could support increases in the system bus speed. Two almost simultaneous system architecture transitions are collapsing the focus of the memory bandwidth packaging challenges to primarily CPU packaging.

The first transition is the transition to an Integrated Memory Controller (IMC) in the CPU. Figure 3 illustrates the shift in the system architecture from that of Figure 1. The system memory now interfaces directly to the CPU through a system memory bus. The entire burden of pin count and interconnect speed to sustain increases in memory bandwidth requirements now falls on the CPU package alone.




Figure 3: System architecture with integrated memory controller on the CPU

The second system architecture transition is the potential move to multi- and many-core CPU architectures. With the increase in cores comes a dramatic transition in memory bandwidth to “feed” the cores, particularly for the class of parallel applications envisioned. Since the

CPU package is now the primary interconnect to the system memory, the CPU package bears the burden of the increase in memory bandwidth. Figures 4 and 5 illustrate the transition in memory bandwidth requirements and the projected CPU package pin count growth, respectively, as the transition to multi- and many-core CPUs with integrated memory controllers occurs.




Figure 4: Bandwidth requirements




Figure 5: Package pin count growth for memory I/Os

One solution to address the bandwidth challenge and to stave off the continued increase in package pin count is to incorporate memory into a CPU + Memory Multichip Package (MCP). This changes the paradigm of packaging and the role of packaging in the memory hierarchy. Whereas packaging has previously been an enabler of higher bandwidths, now packaging would become a crucial sub-level in the overall memory system hierarchy. Moving to on-package memory, however, is not a trivial solution to implement. In the final main section of this

paper, we discuss the challenges and benefits/capabilities of an MCP configuration and various MCP architectures.

First it is beneficial to discuss tera-scale computing challenges in the area of memory bandwidth.

TERA-SCALE COMPUTING MEMORY BANDWIDTH CHALLENGES FOR PACKAGE TECHNOLOGY

Figure 6 shows the historical trend for memory bandwidth demand [3]. Today's bandwidth demand is in the 10–20GB/s range. From Figure 4 it is obvious that the move to multi- and many-core computing will easily drive a need for bandwidth in the 100s GB/s range in the not so distant future. Extrapolating from this, a target of 1TB/s of memory bandwidth for tera-scale computing architectures is not unreasonable [3–5].




Figure 6: Historical trend for memory bandwidth [3]

Re-architecting a system capable of delivering this level of bandwidth is a challenge, given that the traditional methods are already reaching realistic limits. On many microprocessors, SRAM already occupies approximately half of the die real estate [4]. Increasing the amount of on-die memory becomes prohibitive from a cost and die size growth perspective. Increasing the bandwidth of board-level memory becomes prohibitive because of the continued increase in pin count and power with interconnect speed required to sustain bandwidth increases. Using on-package memory becomes a potential attractive intermediary level in the memory hierarchy that can work with chip-level and board-level memory to provide the bandwidth for tera-scale computing applications.

PACKAGE ARCHITECTURES TO MEET THE MEMORY BANDWIDTH NEEDS OF TERA-SCALE COMPUTING

To meet the memory bandwidth needs of tera-scale computing there needs to be an evolutionary transition in package architectures and technologies. In this section we discuss three architectures in detail: the CPU + memory 2D planar MCP, a package substrate embedded memory + CPU MCP architecture, and a 3D CPU + memory stacked die MCP. Figure 7 illustrates these MCP concepts.

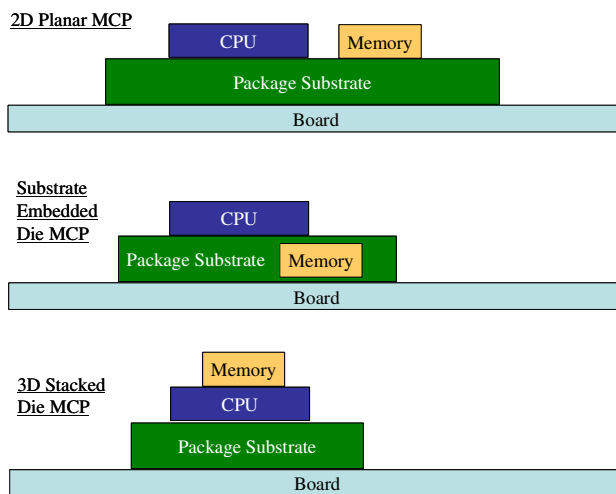


Figure 7: Memory + CPU package architectures for addressing bandwidth challenges

Each of these package architectures has benefits and challenges associated with the technologies. The memory + CPU 2D planar MCP is the most straightforward to implement and can be implemented with today's packaging technologies. There are capability limits on the amount of additional memory bandwidth this architecture can provide, however. Embedding memory in the package is the next evolutionary step in enabling higher memory bandwidth at the package level. This potentially enables higher bandwidths than the memory + CPU 2D planar MCP but requires technology development and comes with integration challenges. The final architecture, 3D stacked die memory + CPU MCP, potentially enables bandwidth capability surpassing the previous two architectures, but requires much technology development and has significant integration challenges. Given these tradeoffs among the architecture challenges and their bandwidth capabilities, a transitional package technology roadmap makes sense and is proposed in Figure 8.

For the remainder of this section we provide details on the capabilities and challenges of each of these on-package memory architectures. For the purposes of this discussion, we assume a memory technology that is capable of

delivering the bandwidth targets in terms of data rate and connection density that will be discussed. A discussion of the memory technology details is outside the scope of this paper. Also, the focus of our discussion is on memory bandwidth. Memory capacity and latency also play a key role in CPU system performance, but are outside the scope of this paper.

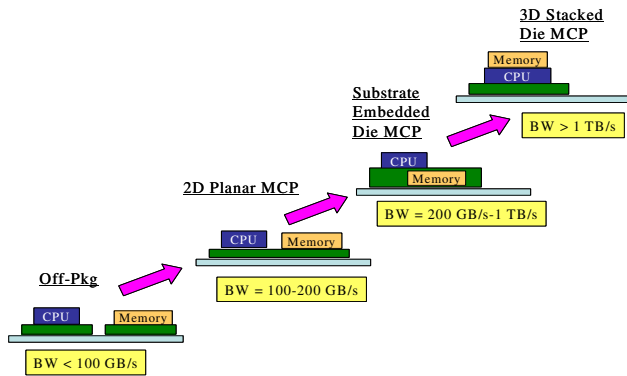


Figure 8: Proposed roadmap of package architecture transitions to address the memory bandwidth challenge

The first transition to on-package memory that can be implemented in the most evolutionary manner with respect to today's package technology is the memory + CPU 2D planar MCP. Intel has used 2D planar MCP packaging in the past and continues to use MCP package technology today for many of its multi-core processors, so this package technology is not a revolutionary technology. There are unique challenges associated with a CPU + memory 2D planar MCP that realistically limit its bandwidth capabilities.

The key challenges with a CPU + memory 2D planar MCP are that heterogeneous die are being assembled onto a single-package substrate with a requirement of optimizing performance to achieve a very fast, dense memory bus interconnection scheme. There are both design and electrical performance challenges associated with this architecture. Key design challenges are form factor fit, die placement, and routability.

In general, bump pitch between the CPU die and the memory or DRAM die will likely not match. This leads to routing issues that do not enable the design to take full advantage of line/space density capabilities of the package technology. Figure 9 illustrates a typical scenario when trying to route 200 I/Os between two die on an MCP substrate. Single-layer routing becomes impossible because of cut-off of the routing lanes. The solution is to use two-layer routing. This results in an increase in package cost and challenges in performance resulting from the division of the memory bus into two layers of routing.




Figure 9: Design and routing challenges with on-package memory 2D planar MCP

In addition to routing challenges, there are challenges in large numbers of I/Os escaping from each die due to bump pitch constraints. Table 1 summarizes results for 200 and 400 I/Os. While routing a memory bus with 200 I/Os is fairly scalable with reasonable package technology and bump pitch routing capability scaling assumptions, increasing this to 400 I/Os becomes challenging.

Table 1: Design and routing/die-escape challenges with on-package memory 2D planar MCP

Attributes / Pitch	200 I/Os		400 I/Os	
	150 μm	175 μm	150 μm	175 μm
Bump Pitch	150 μm	175 μm	150 μm	175 μm
Bumps / Row	72	62	72	62
# of Rows Needed (on-die)	3	4	6	7
I/O Region Depth (on-die)	0.5 mm	0.7 mm	0.9 mm	1.1 mm
Max. # of Lines Escapable thru Min. Pitch	2	3	2	3
Max. # of Rows Escapable	4	5	4	5
Pkg Layers Needed for Die Escape	1	1	2	2

Signal integrity issues in an MCP configuration also lead to performance challenges. Because there is still a substantial amount of trace length that is routed on the package between chips and these are routed very densely, crosstalk limits performance. For a single-ended configuration, the upper limit is ~6–7 Gb/s. Results of a signal integrity sensitivity study are summarized in Figures 10 and 11.




Figure 10: 2D planar MCP with on-package memory signal integrity results (eye height)




Figure 11: 2D planar MCP with on-package memory signal integrity results (eye width)

Combining the limits introduced by routing, fit, and signal integrity challenges, an estimate on the maximum sustainable bandwidth of a 2D planar MCP configuration is 100GB/s–200GB/s, depending upon the number of memory chips placed on the MCP. This also assumes a transition in memory technology to enable the types of connection densities and memory speeds used in this study. While this is a substantial amount of memory bandwidth capability, it is still not sufficient to meet the ultimate targets for tera-scale computing in the long term.

The next transition in package technology that can enable higher memory bandwidth than the CPU + memory MCP is a package embedded memory architecture. Figure 12 shows a schematic of this package architecture. This type of package architecture eliminates one level of transition between the CPU die and the DRAM die, i.e., the routing, which is responsible for the crosstalk that limits the ultimate I/O speed of the CPU + memory MCP architecture. This conceivably enables higher bandwidth by providing a very short and direct CPU-to-DRAM interconnect.




Figure 12: MCP with substrate embedded DRAM

Signal integrity simulations for a typical CPU-to-DRAM interconnect using a substrate embedded DRAM revealed very minimal impact due to crosstalk since the die-to-die connections are separated by an appreciable distance, equal to at least the bump pitch. The model used for the signal integrity studies is shown in Figure 13. At 4Gb/s, the substrate embedded memory configuration resulted in at least 30–40% more margin than the CPU + memory 2D planar MCP configuration results shown in Figure 10. Extrapolating from these results, it is conceivable that the substrate embedded memory architecture can easily achieve a bit rate of at least 10Gb/s. Given that direct connections between the CPU and memory can be made at the same pitch as the die bump pitch, hundreds of connections can be enabled in a small area. Consequently, this architecture can enable a memory bandwidth in the 200GB/s–1TB/s range.




Figure 13: MCP with substrate embedded DRAM

One challenge with this type of architecture is the thermal performance with an embedded die. Preliminary thermal modeling and historical data suggest that a limit of

approximately ten watts or less for the embedded device should be maintained to avoid excessive refresh rates and increased power penalties. There are also substantial integration challenges with this architecture. This is, however, an intriguing architecture for enabling bandwidths in the 200GB/s-1TB/s range.

To enable memory bandwidths beyond 1TB/s, the 3D CPU + memory stacked die MCP architecture becomes interesting. Because this will provide the shortest possible interconnect between the CPU and memory die, the bit rate will far exceed 10Gb/s. In addition, the interconnect density will scale to enable thousands of die-to-die interconnects. Intel recently demonstrated a single-chip teraflop processor, Polaris, with 80 cores. Polaris contains hooks for stacking a separate SRAM chip, Freya, in a 3D configuration [6] and [7].

SUMMARY AND CONCLUSIONS

Memory bandwidth is one of the key challenges associated with tera-scale computing. Package technology will play a key role in answering that challenge. In this paper, we reviewed the historical trends in memory bandwidth and their impact on package technology. Key challenges in the past have been a continuing pin count growth that leads to package body size growth as well as design challenges. In addition, there are fundamental limits to the off-package memory bus speed that can be supported. A transition to on-package memory is necessary for supporting tera-scale computing needs. We reviewed a transitional, evolutionary roadmap for package technology to implement on-package memory architectures capable of meeting the needs of tera-scale

computing. Table 2 summarizes the technology and architecture options, features, capabilities, and limits of each.

Table 2: Summary of features and capabilities of on-package memory architectures

	On-Die	On-Package			Off-Pkg
		2D Planar MCP	Substrate Embedded Die MCP	3D Stacked Die MCP	
I/O Speed		4 – 7 Gb/s	4 – 10 Gb/s	10+ Gb/s	4 Gb/s
# of Connections		200 – 400	400 – 10,000	10,000+	~4 – 6 channels DDR
Bandwidth		100 – 200 GB/s	200 GB/s – 1 TB/s	> 1 TB/s	< 100 GB/s
Capacity	< 50 MB	512+ MB	< 512 MB	< 512 MB	> 1 GB
Interconnect Length	< 1 mm	~7 mm	1-3 mm	1 mm	> 100 mm

ACKNOWLEDGMENTS

The authors thank their colleagues who have worked with them on many aspects of this work, including Gans Ganesan, Tom Dory, Bhanu Jaiswal, Jemmy Sutanto, Jiangqi He, Anne Augustine, Ji Wu, Kinya Ichikawa, Yoshihiro Tomita, Sriram Dattaguru, Yonggang Li, Haluk Balkan, Ravi Mahajan, Debendra Mallik, Je-Young Chang, Greg Chrysler, Sriram Muthukumar, Jim Irvine, Islam Salama, Ihtesham Chowdhury, Aleksander Aleksov, and Henning Braunisch. We also acknowledge the continuing support of our managers: Kaladhar Radhakrishnan, Gaurang Choksi, Huong Do, Ken Brown, Bob Sankman, and Kemal Aygun. We thank the reviewers of this paper for their constructive input. Finally, we thank Jerry Bautista for serving as our mentor during the writing of this paper.

REFERENCES

- [1] Held, J., Bautista, J. and Koehl, S., “From a Few Cores to Many: A Tera-scale Computing Research Overview,” *Research at Intel White Paper*, download.intel.com/research/platform/terascale/terascale_overview_paper.pdf, 2006.
- [2] Hsu, L., Iyer, R., Makineni, S., Reinhardt, S. and Newell, D., “Exploring the Cache Design Space for Large Scale CMPs,” *ACM SIGARCH Computer Architecture News*, vol. 33, no. 4, Sept. 2005, pp. 24–33.
- [3] Mallik, D. Radhakrishnan, K., He, J., Chiu, C., Kamgaing, T., Searls, D. and Jackson, J. D., “Advanced Package Technologies for High-Performance Systems,” *Intel Technology Journal*, vol. 9, no. 4, Nov. 9, 2005, pp. 259–271.
- [4] Moore, S. K., “Winner: Masters of Memory,” *IEEE Spectrum*, Jan. 2007, pp. 45–49.

- [5] Patti, R., "Design and Application of 3D Memories," Proceedings of IMAPS International Conference and Exhibition on Device Packaging, Feb. 27, 2007.
- [6] Vangal, S., et al., "An 80-Tile 1.28TFLOPS Network-on-Chip in 65 nm CMOS," in *Proceedings of ISSCC 2007 (IEEE International Solid-State Circuits Conference)*, Feb. 12, 2007.
- [7] Johnson, R.C., "Intel's Teraflops Chip Uses Mesh Architecture to Emulate Mainframe," *EE Times Online*, at www.eetimes.com/showArticle.jhtml?articleID=197004697, Feb. 12, 2007.

AUTHORS' BIOGRAPHIES

Lesley Anne Polka joined Intel in 1994. She is a Staff Electrical Packaging Engineer and works in the Assembly and Test Technology Development Division in Chandler, Arizona. Her focus is electrical issues related to development of Intel's packaging technologies for CPU, chipset, and wireless/RF packaging. During her twelve and a half years at Intel she has worked on various aspects of electrical packaging, including high-speed I/O and

signal integrity, power delivery, analysis and characterization. She is presently working on new CPU package architectures to meet memory bandwidth and power-delivery challenges for future Intel microprocessors. Lesley has published over 15 papers on topics related to electrical design and technology development for microprocessor packaging and computational electromagnetics. She has two U.S. patents in packaging and a third U.S. patent pending in this area. Lesley obtained her B.S., M.S. and Ph.D. degrees, all in Electrical Engineering, from Arizona State University, in 1987, 1989 and 1995, respectively. Her e-mail is lesley.a.polka@intel.com.

Huthasana (Huthas) Kalyanam has been with Intel Corporation for the past one and a half years. He works as a Packaging Engineer within the Design Group in the Assembly and Test Technology Development Division in Chandler, Arizona. His primary job responsibilities include designing packages that are optimized for cost, electrical performance, and form factor. He designs packages that cater to Intel's Server and Desktop market segments. His responsibilities also include designing test packages for Intel's next-generation path-finding efforts in Packaging and Silicon. He is currently working on designing packages for multichip designs. Huthas received his M.S. degree in Electrical Engineering with a focus on VLSI and Analog systems from the University of Colorado at Boulder in 2004, after which he worked for a year with Instec Inc. Boulder as a Product and Design

Engineer for Instec's Embedded Systems. His e-mail is huthasana.kalyanam@intel.com.

Grace Hu joined Intel in 2005 as an Electrical Packaging Engineer. She works on high-speed I/O characterization for CPU and chipset packages. Her focus areas also include high-frequency measurement and metrology development, broad-band dielectric material characterization, and risk assessments for new package technologies. Grace obtained her B.S. degree from Shanghai Jiao Tong University and her M.S. degree from the University of Missouri-Rolla, both in Electrical Engineering in 1999 and 2002, respectively. Her e-mail is grace.hu@intel.com.

Satish Krishnamoorthy joined Intel in 2005 and works as an Electronic Packaging Engineer in the Assembly and Test Technology Development Division at Chandler, Arizona. His focus is in package electrical analysis and design with particular emphasis on power delivery and signal integrity for various CPU and chipset products. Satish has published over five papers in refereed journals in the area of electronic circuits. He obtained his B.S. degree in Electronics and Communication Engineering from Anna University and his M.S. degree in Electrical Engineering from Arizona State University, in 2003 and 2005, respectively. His e-mail is satish.krishnamoorthy@intel.com.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel's trademarks may be used publicly with permission only from Intel. Fair use of Intel's trademarks in advertising and promotion of Intel products requires proper acknowledgement.

*Other names and brands may be claimed as the property of others.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS® Ready mark under license from Palm, Inc.

Copyright © 2007 Intel Corporation. All rights reserved.

This publication was downloaded from
<http://www.intel.com>.

Additional legal notices at:
<http://www.intel.com/sites/corporate/tradmarx.htm>.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

developer.intel.com/technology/itj/index.htm